

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re PATENT APPLICATION of

Steven Robert HETZLER)		
)	Group:	2133
Serial No.: 10/619,649)		
)	Examiner:	M.M. Chaudry
Filed: July 14, 2003)		
)	Atty Dkt:	ARC920030015US1
For: AUTONOMIC PARITY)		
EXCHANGE)	Conf. No.	8646

AMENDMENT AFTER EX PARTE QUAYLE ACTION

MAIL STOP AMENDMENT

Commissioner For Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450

Sir:

In response to the Office Action (Confirmation No. 8646) mailed May 11, 2006, please amend the above-captioned application as set forth below, and reconsider this application in view of the following remarks.

Amendments to the Specification begin on page 2 of this paper.

Amendments to the Claims are reflected in the listing of the claims that begins on page 6 of this paper.

Amendments to the Drawings begin on page 16 of this paper and include both an attached replacement sheet and an annotated sheet showing changes.

Remarks begin on page 17 of this paper.

An Appendix including amended drawing sheets is attached following page 19 of this paper.

AMENDMENTS TO THE SPECIFICATION:

Please replace the heading appearing prior to Paragraph 1 with the following heading:

~~CROSS-REFERENCE~~CROSS-REFERENCED TO RELATED APPLICATIONS

Please replace the disclosure contained in Paragraph 1 with the following rewritten disclosure:

The present application is related to Patent Application Serial No. 10/619,641~~(Attorney Docket No. ARC9-2003-0014-US1)~~, entitled "Anamorphic Codes", Patent Application Serial No. 10/619,633~~(Attorney Docket No. ARC9-2003-0016-US1)~~, entitled "Multi-path Data Retrieval From Redundant Array," and Patent Application Serial No. 10/619,648~~(Attorney Docket No. ARC9-2003-0040-US1)~~, entitled "RAID 3 + 3" each co-pending, co-assigned and filed concurrently herewith, and each incorporated by reference herein. The present application is also related to co-pending and co-assigned Patent Application Serial No. 10/600,593~~(Attorney Docket No. YOR9-2003-0069-US1)~~, which is also incorporated by reference herein.

Please replace the disclosure contained in Paragraph 22 with the following rewritten disclosure:

The advantages of the present invention are provided by a method for increasing failure tolerance of a storage system having a plurality of arrays such that each array has a plurality of storage units. The arrays of the storage system include redundancy based on an erasure or error correcting code, such as a parity code, a Winograd code, a symmetric code, a Reed-Solomon code, an EVENODD code or a derivative of an EVENODD code. The failure tolerance of a storage system is given by the minimum Hamming distance D of the system. The minimum Hamming distance of the system is, accordingly, the minimum of all the minimum Hamming distances of the respective stripes, that is, $D = \min(d_i)$. A donor array is selected from the plurality of arrays when the difference between a minimum Hamming distance of the donor array and a minimum Hamming distance of a recipient array is greater or equal to 2. A donor storage unit is selected in the donor array based on a minimal performance impact on the donor array. A

recipient storage unit is selected from the recipient array. At least a portion of lost information is then rebuilt from the recipient array onto the selected storage unit in the donor array. The recipient information is selected based on an improved performance of the recipient array. The selected storage unit is indicated to the donor array as having been donated before the lost information is rebuilt on the selected storage unit. Preferably, the minimum Hamming distance of the recipient array is $d \geq 2$ before the donor array is selected from the plurality of arrays. When a spare storage unit becomes available, the spare storage unit is assigned to a selected array in a conventional manner.

Please replace the disclosure contained in Paragraph 25 with the following rewritten disclosure:

Further still, the present invention provides a data storage system having a plurality of arrays and a system array controller. Each array has a plurality of storage units and includes redundancy based on an erasure or error correcting code, such as a parity code, a Winograd code, a symmetric code, a Reed-Solomon code, an EVENODD code or a derivative of an EVENODD code. The system array controller is coupled to each array and detects a failure of a storage unit in a first array of the plurality of arrays. The system controller then selects a storage unit in a second array of the plurality of arrays when a difference between a minimum Hamming distance of the second array and a minimum Hamming distance of the first array is greater or equal to 2. At least a portion of lost information from the first array is rebuilt onto the selected storage unit of the second array. Each storage unit can be an HDD, a volatile Random Access Memory device, a non-volatile Random Access Memory device, an optical storage device, or a tape storage device.

Please replace the disclosure contained in Paragraph 40 with the following rewritten disclosure:

APX allows arrays within a set of arrays to exchange redundancy, thereby overcoming exposure to failures that are concentrated on a single array of the set of arrays. For example, if a first array

has a minimum Hamming distance that is less than the minimum Hamming distance of a second array by 2 or more, the second array can donate a storage unit to the first array. Afterward, the failure tolerance of the first array will be increased and the failure tolerance of the second array will be reduced, but to a level that is not less than the first array. Accordingly, the minimum Hamming distance of the system will be increased, thereby increasing the failure tolerance of the system.

Please replace the disclosure contained in Paragraph 42 with the following rewritten disclosure:

Figure 1b shows an exemplary storage system, indicated generally as 150, comprising two storage arrays 153 and 154, each comprising multiple storage units 155, that are respectively connected to different array controllers 152 and 151. Storage array 153 communicates with array controller 152 over interface 157, and storage array 154 communicates with array controller 151 over interface 156. Array controllers 151 and 152 respectively communicate with other array controllers and storage systems over interfaces 158 and 159. Also shown in Figure 1b is a communication connection 160 that allows array controllers 151 and 152 to communicate with each other.

Please replace the disclosure contained in Paragraph 59 with the following rewritten disclosure:

APX can be used beyond simply increasing the minimum Hamming distance of a storage system. Many other factors may be included in determining whether to perform APX and to choose donors and recipients. For example, the individual failure probabilities of components when they are non-uniform, the combinations of failures that lead to data loss, and the effects on system performance may all be considered. In such cases, the minimum Hamming distance of the system could remain unchanged following APX.

Please replace the disclosure contained in Paragraph 60 with the following rewritten disclosure:

APX can be used with other array types having minimum Hamming distance $d \geq 3$. Additionally, a smaller array size allows APX to be used more efficiently, and allows large systems consisting of small arrays to achieve high failure tolerance. When a storage system has a spare pool, it is best to perform rebuilds onto the spare pool before performing an APX operation.

Please replace the Abstract of the Disclosure with the following rewritten Abstract of the Disclosure:

Error tolerance is increased for a storage system having a plurality of arrays by making local redundancy in a selected array globally available throughout the storage system. To achieve the increased error tolerance, a donor array is selected from the plurality of arrays when the difference between a minimum Hamming distance of the donor array and a minimum Hamming distance of a recipient array is greater or equal to 2. A donor storage unit is selected in the donor array and recipient information is then rebuilt from the recipient array on the selected storage unit. The selected storage unit is indicated to the donor array as having been donated before the lost information is rebuilt on the selected storage unit. Preferably, the minimum Hamming distance of the recipient array is $d \geq 2$ before the donor array is selected from the plurality of arrays.

AMENDMENTS TO THE CLAIMS:

The following listing of the claims replaces all prior versions and listings of the claims in the present application:

Please cancel claims 8, 9, 26-44 and 66-87 without prejudice or disclaimer of the subject matter contained therein.

Please amend claims 1, 24, 45 and 48 as follows:

1. (currently amended) A method for increasing failure tolerance of a storage system having a plurality of arrays, each array having a plurality of storage units, the method comprising steps of:

selecting a recipient array from the plurality of arrays;
selecting a donor array from the plurality of arrays when a difference between a minimum Hamming distance of the donor array and a minimum Hamming distance of the recipient array is greater or equal to 2;
selecting a donor storage unit in the donor array; and
rebuilding at least a portion of lost recipient information from the recipient array on the selected storage unit in the donor array.

2. (original) The method according to claim 1, wherein the minimum Hamming distance of the recipient array is $d \geq 2$ before the step of selecting the donor array from the plurality of arrays.

3. (original) The method according to claim 1, further comprising a step of indicating to the donor array that the selected storage has been donated before the step of rebuilding the lost information on the selected storage unit.

4. (original) The method according to claim 1, wherein the storage units are hard disk drives.

5. (original) The method according to claim 1, wherein the storage units are RAM storage devices.

6. (original) The method according to claim 1, further including a step of selecting a recipient storage unit from the recipient array.

7. (original) The method according to claim 6, wherein at least a portion of the lost contents of the recipient storage unit are rebuilt onto the donor storage unit.

8. (canceled)

9. (canceled)

10. (original) The method according to claim 1, wherein the arrays of the storage system include redundancy based on an erasure or error correcting code.

11. (original) The method according to claim 10, wherein the erasure or error correcting code is a parity code.

12. (original) The method according to claim 10, wherein the erasure or error correcting code is a Winograd code.

13. (original) The method according to claim 10, wherein the erasure or error correcting code is a symmetric code.

14. (original) The method according to claim 10, wherein the erasure or error correcting code is a Reed-Solomon code.

15. (original) The method according to claim 10, wherein the erasure or error correcting code is an EVENODD code.

16. (original) The method according to claim 10, wherein the erasure or error correcting code is a derivative of an EVENODD code.

17. (original) The method according to claim 10, wherein the arrays of the storage system includes redundancy based on a product of a plurality of erasure or error correcting codes.

18. (original) The method according to claim 17, wherein at least one of the erasure or error correcting codes is a parity code.

19. (original) The method according to claim 17, wherein at least one of the erasure or error correcting codes is a Winograd code.

20. (original) The method according to claim 17, wherein at least one of the erasure or error correcting code is a symmetric code.

21. (original) The method according to claim 17, wherein at least one of the erasure or error correcting code is a Reed-Solomon code.

22. (original) The method according to claim 17, wherein at least one of the erasure or error correcting code is an EVENODD code.

23. (original) The method according to claim 17, wherein at least one of the erasure or error correcting code is a derivative of an EVENODD code.

24. (currently amended) The method according to claim 1, wherein when a storage unit in the donor array fails during the step of rebuilding at least a portion of recipient information from the recipient array on the selected storage unit, the method further comprising steps of:

terminating the step of rebuilding at least a portion of recipient information from the recipient array on the selected storage unit;

selecting a second donor array from the plurality of arrays when a difference between a minimum Hamming distance of the second donor array and a minimum Hamming distance of the second recipient array is greater or equal to 2;

selecting a donor storage unit in the second donor array; and

rebuilding at least a portion of lost recipient information from the recipient array on the selected storage unit in the second donor array.

25. (original) The method according to claim 1, wherein when a spare storage unit becomes available, the method further comprising a step of assigning the spare storage unit to a selected array.

26. (canceled)

27. (canceled)

28. (canceled)

29. (canceled)

30. (canceled)

- 31. (canceled)
- 32. (canceled)
- 33. (canceled)
- 34. (canceled)
- 35. (canceled)
- 36. (canceled)
- 37. (canceled)
- 38. (canceled)
- 39. (canceled)
- 40. (canceled)
- 41. (canceled)
- 42. (canceled)
- 43. (canceled)
- 44. (canceled)

45. (currently amended) A data storage system, comprising:
a plurality of arrays, each array having a plurality of storage units; and
a system array controller coupled to each array, the system array controller detecting a failure of a storage unit in a first array of the plurality of arrays, selecting a storage unit in a second array of the plurality of arrays when a difference between a minimum Hamming distance of the second array and a minimum Hamming distance of the first array is greater or equal to 2, and rebuilding at least a portion of information from the first array onto the selected storage unit of the second array.

46. (original) The data storage system according to claim 45, wherein when a spare unit becomes available, the spare unit is assigned to the second array.

47. (original) The data storage system according to claim 45, wherein at least one array has a non-uniform failure probability.

48. (currently amended) The data storage system according to claim 45, wherein the system array controller includes a plurality of array controllers, an array controller being coupled to at least one array of the plurality of arrays, each respective array controller detecting a failure of a storage unit in each array associated with the array controller, a first array controller of the plurality of array controllers selecting a storage unit in an array associated with the first array controller when a difference between a minimum Hamming distance of the array of the selected storage unit and a minimum Hamming distance of an array associated with a second array controller of the plurality of array controllers is greater or equal to 2, and the first and second array controllers rebuilding at least a portion of lost information from the array associated with the second array controller onto the selected storage unit in the array associated with the first array controller.

49. (original) The data storage system according to claim 48, wherein when a spare

unit becomes available, the spare unit is assigned to the array of the selected storage unit.

50. (original) The data storage system according to claim 48, wherein at least one array has a non-uniform failure probability.

51. (original) The data storage system according to claim 45, wherein the arrays of the data storage system include redundancy based on an erasure or error correcting code.

52. (original) The data storage system according to claim 51, wherein the erasure or error correcting code is a parity code.

53. (original) The data storage system according to claim 51, wherein the erasure or error correcting code is a Winograd code.

54. (original) The data storage system according to claim 51, wherein the erasure or error correcting code is a symmetric code.

55. (original) The data storage system according to claim 51, wherein the erasure or error correcting code is a Reed-Solomon code.

56. (original) The data storage system according to claim 51, wherein the erasure or error correcting code is an EVENODD code.

57. (original) The data storage according to claim 51, wherein the erasure or error correcting code is a derivative of an EVENODD code.

58. (original) The data storage system according to claim 45, wherein the arrays of the storage system includes redundancy based on a product of a plurality of erasure or error

correcting codes.

59. (original) The data storage system according to claim 58, wherein at least one of the erasure or error correcting codes is a parity code.

60. (original) The data storage system according to claim 58, wherein at least one of the erasure or error correcting codes is a Winograd code.

61. (original) The data storage system according to claim 58, wherein at least one of the erasure or error correcting codes is a symmetric code.

62. (original) The data storage system according to claim 58, wherein at least one of the erasure or error correcting codes is a Reed-Solomon code.

63. (original) The data storage system according to claim 58, wherein at least one of the erasure or error correcting codes is an EVENODD code.

64. (original) The data storage system according to claim 58, wherein at least one of the erasure or error correcting codes is a derivative of an EVENODD code.

65. (original) The data storage system according to claim 45, wherein each storage unit is one of a Hard Disk Drive, a volatile Random Access Memory device, a non-volatile Random Access Memory device, an optical storage device, and a tape storage device.

66. (canceled)

67. (canceled)

- 68. (canceled)
- 69. (canceled)
- 70. (canceled)
- 71. (canceled)
- 72. (canceled)
- 73. (canceled)
- 74. (canceled)
- 75. (canceled)
- 76. (canceled)
- 77. (canceled)
- 78. (canceled)
- 79. (canceled)
- 80. (canceled)
- 81. (canceled)

82. (canceled)

83. (canceled)

84. (canceled)

85. (canceled)

86. (canceled)

87. (canceled)

AMENDMENT TO THE DRAWINGS

The attached sheets of drawings include a replacement sheet and an annotated sheet for Figures 1a and 1b. In Figure 1a, labels have been added reference numerals 102, 103 and 104 and in Figure 1b, labels have been added to reference numerals 153, 154 and 155. Additionally, in Figure 1b, interface 157 has been connected to controller 152.

Attachment: Replacement Sheet For Figures 1a and 1b (1 sheet)

Annotated Sheet Showing Changes for Figures 1a and 1b (1 sheet)

REMARKS

The Office Action dated May 11, 2006, has been carefully reviewed, and in view of the above changes, the present application should now be ready for to be passed to issue.

The Amendments To The Claims

Applicants have amended claims 1, 24, 25 and 48 to improve their respective forms in accordance with U.S. patent law. In particular, claims 1, 24, 25 and 48 have been amended by changing the term “minimum distance” to “minimum Hamming distance, and claims 8, 9, 26-44 and 66-87 have been canceled. Additionally, claim 1 has been amended to further improve its form in accordance with U.S. patent law by requiring that a recipient array is selected “from the plurality of arrays.”

The Objection To The Drawings

The drawings stand objected to because in Figure 1a reference numerals 102, 103 and 104 and in Figure 1a reference numerals 153, 154 and 155 need labels.

Applicants have attached a replacement sheet and an annotated sheet for Figures 1a and 1b. In Figure 1a, labels have been added reference numerals 102, 103 and 104 and in Figure 1b, labels have been added to reference numerals 153, 154 and 155. Additionally, in Figure 1b, interface 157 has been connected to controller 152. Support for this amendment to Figure 1b can be found throughout the specification, for example, in Paragraph 42.

Consequently, Applicants respectfully request that the Examiner withdraw this objection.

The Objection To The Disclosure

The disclosure stands objected to for various informalities relating to the term “minimum distance,” and relating to the identification of co-pending patent applications.

Applicants have amended the specification, as set forth above, to change the term “minimum distance” to now be “minimum Hamming distance,” and to identify co-pending patent applications by their respective serial numbers.

Additionally, Applicants have amended Paragraph 42 of the disclosure in order to describe reference numeral 155, which is shown in Figure 1b. Support for the added disclosure to Paragraph 42 can be found in Figures 1a and 1b.

Consequently, Applicants respectfully request that the Examiner withdraw this objection.

The Objection To The Claims

Claims 1 and 45 stand objected to for various informalities.

Applicants have amended claims 1 and 45 by changing the term “minimum distance” to “minimum Hamming distance. Additionally, claims 24 and 48 have been similarly amended, and claim 1 has been amended by requiring that a recipient array is selected “from the plurality of arrays.”

Consequently, Applicants respectfully request that the Examiner withdraw this objection.

The Rejection Under 35 U.S.C. § 112, Second Paragraph

Claims 8 and 9 stand rejected under 35 U.S.C. § 112, second paragraph, as indefinite for failing to particularly point out and distinctly claim the subject matter regarded as the invention.

Applicants have canceled claims 8 and 9.

Consequently, Applicants respectfully request that the Examiner withdraw this rejection.

CONCLUSION


In view of the above amendments that place the present application in condition to be passed to issue, should the Examiner find that a telephonic or personal interview would further expedite passage to issue of the present application, the Examiner is encouraged to contact the undersigned attorney at the telephone number indicated below.

A general authorization under 37 C.F.R. § 1.25(b), second sentence, is hereby given to credit or debit Deposit Account No. 09-0441 for the instant filing and for any other fees during the pendency of this application under 37 C.F.R. §§ 1.16, 1.17 and 1.18.

It is requested that this application be passed to issue with claims 1-7, 10-25 and 45-65.

Respectfully submitted,

Date: May 26, 2006

A handwritten signature in black ink, appearing to read "J. Curtin", is written over a horizontal line.

Joseph P. Curtin
Registration No. 34,571
1469 N.W. Morgan Lane
Portland, Oregon 97229-5291
(503) 296-8373
(503) 297-0452 (facsimile)

REPLACEMENT




FIG. 1a




FIG. 1b

ANNOTATED




FIG. 1A




FIG. 1B